IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant:	Todd R. Burkey	Examiner:	Yaima Campos
Serial No.:	10/629,415	Group Art Unit:	2185
Filing Date:	July 29, 2003	Docket No.:	3916
Title	Method, Apparatus and Program Storage Device for Dynamically Resizing Mirrored Virtual Disks in a RAID Storage System		

BRIEF ON APPEAL

Mail Stop Appeal Brief - Patents
Commissioner for Patents
Alexandria, VA 22313

Sir:

  This Brief on Appeal is filed pursuant to the Notice of Appeal filed on October 19, 2010 and is an appeal from the Office Action mailed from the U.S. Patent and Trademark Office on May 21, 2010. The balance of this appeal is set forth under appropriate headings, as specified by 37 C.F.R. 41.37(c).

TABLE OF CONTENTS

I. REAL PARTY IN INTEREST

The real party in interest is XIOtech Corporation, 6455 Flying Cloud Drive, Eden Prairie, MN, 55344, the assignee of the entire right, title and interest in the subject application, by virtue of an assignment recorded on July 29, 2003 at Reel 014353, Frame 0214.

II. RELATED APPEALS AND INTERFERENCES

Appellant, the undersigned Attorney, and Assignee are not aware of any related appeals or interferences that will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

III. STATUS OF CLAIMS

Claims 1-5, 7, 9, 10, 13, and 16-24 are currently pending in this application, and stand finally rejected. Specifically, claims 4 and 16 stand rejected under Section 112, and all the pending claims stand rejected under Section 103. Claims 6, 8, 11, 12, 14, and 15 were canceled in an amendment filed April 28, 2008. Claims 1-3, 9, and 16-24 are being appealed; copies of these claims appear in the Claims Appendix of this Brief.

IV. STATUS OF AMENDMENTS

No amendment was filed subsequent to final rejection.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The subject matter of each of the independent claims involved in the appeal is summarized below. In the case of claim 16, every means plus function is identified, and the structure, material, or acts corresponding to each claimed function is set forth with reference to the specification by page and line number, and to the drawing, if any, by reference character.

A. Claim 1

Claim 1 recites a program storage device readable by a computer embodying in a tangible medium one or more programs of instructions executable by the computer to perform a method for dynamically expanding mirrored virtual disks in a virtual disk storage system, the method comprising:

> receiving by a source virtual disk a request to dynamically expand the mirrored virtual disks, which include the source virtual disk and at least one destination virtual disk;

associating additional storage with the mirrored virtual disks;

increasing respective sizes of each of the at least one destination virtual disk before reporting a new storage size of the source virtual disk; and

reporting the new size of the source virtual disk.

The application teaches methods for dynamically expanding and contracting a set of virtual disks, including a source virtual disks and at least one destination virtual disk that mirrors the source. Claim 1 is concerned with the dynamic expansion capability. During expansion, the increasing step precedes the reporting step to ensure that all mirroring disks have the desired capacity to accommodate any new data written to the set of mirroring disks before such writing occurs.

Aspects of the above claim are supported at least by the following portions of the specification and drawings:

- program storage device, instructions, computer: Fig.2 refs. 268, 290, and 267; page 7 lines 9-15; page 7 line 16-page 8 line 2; page 14 line 18-page 15 line 5.
- receiving request by source virtual disk: Fig. 4 ref. 410; page 7 lines 9-15.
- associating additional storage with the mirrored virtual disks: page 12 lines 7-20; page 15 line 14-page 16 line 2; Fig. 3; Fig. 5 ref. 510 and 520.
- increasing sizes before reporting: page 12 lines 7-20; page 15 line 14-page 16 line 2; page 16 line 17-page 17 line 7; Fig. 5 ref. 530.

B. Claim 20

Claim 20 is a claim for a method, comprising:

> receiving a request to dynamically resize mirrored virtual disks, the mirrored virtual disks comprising a source virtual disk and a set of destination virtual disks that includes at least one destination virtual disk;
> associating additional storage with the mirrored virtual disks;
> increasing respective new storage sizes of each destination virtual disk before reporting a new storage size of the source virtual disk; and
> reporting the new storage size of the source virtual disk.

Claim 20 is a claim related to dynamic resizing by expansion. The relevant references are a subset of those supporting claim 1, the subset being repeated here for completeness:

Aspects of the above claim are supported at least by the following portions of the specification and drawings:

- receiving request by source virtual disk: Fig. 4 ref. 410; page 7 lines 9-15.

- associating additional storage with the mirrored virtual disks: page 12 lines 7-20; page 15 line 14-page 16 line 2; Fig. 3; Fig. 5 ref. 510 and 520.
- increasing sizes before reporting: page 12 lines 7-20; page 15 line 14-page 16 line 2; page 16 line 17-page 17 line 7; Fig. 5 ref. 530.

C. Claim 23

Claim 23 is a claim for an apparatus, comprising:

> a set of mirrored virtual disks, including a source virtual disk and at least one destination virtual disk, the at least one destination virtual disk mirroring the source virtual disk, wherein the source and destination virtual disks have the same size;
>
> a management module that includes
>
>> a host side interface adapted to communicating with host devices, through which the management module is adapted by logic to report the size of the mirrored virtual disks and to receive a request to expand the mirrored virtual disks, and
>>
>> a storage system interface for communicating with the virtual disks that is adapted to requesting the source virtual disk to expand and to obtain reports of the size of the virtual disks from the source virtual disk; and
>
> logic adapted to
>
>> provide reports of the size of the source virtual disk to the management module through the storage system interface,
>>
>> satisfy an expansion request by creating an amount of necessary storage before changing the size that will be obtained by the management module in reports from the source virtual disk; and

> change the size that will be obtained by the management module
> in reports from the source virtual disk.

Aspects of the above claim, which, like claim 1, also pertains to expansion, are supported at least by the following portions of the specification and drawings:

- Set of mirrored virtual disks (page 15 line 14-page 16 line 2), including at least a source and a destination (page 12 lines 7-20) that have the same size (page 15 line 14-page 16 line 2);
- Management module with host and storage system interfaces: Fig. 2;
- Host side interface: Fig. 1 ref. 120, page 13 lines 1-5 and page 13 lines 6-9 (network, Ethernet); Fig. 1 ref. 122 and page 13 lines 6-9 (storage area network); Fig. 1 ref. 124 and page 13 lines 6-9 (point-to-point connection); 140 and page 13 lines 10-14 (network connection); and Fig. 1 ref. 134 and (virtual representation); Fig. 2 ref. 210 and page 14 lines 1-7 (bus interface); 214 and page 14 lines 8-13 (Fibre Channel Arbitrated Loop); Fig. 2 ref. 216 and page 14 lines 1-7 (point-to-point connection), Fig. 2 ref. 222 and page 14 lines 8-13 (Fibre Channel), switched fabric page 14 lines 1-7.
- Storage system interface: Fig. 1 ref. 122 and page 13 lines 6-9 (storage area network); 124 and page 13 lines 6-9 (point-to-point connection); Fig. 1 ref. 140 and page 13 lines 10-14 (network connection), and Fig. 1 ref. 134 and page 13 lines 15-20 (virtual representation of disk); Fig. 2 ref. 200 and 240 and page 13 lines 10-14 (management module connected to storage).
- Logic: Fig. 2 ref. 267, 292, 268, 290; Fig. 5; page 14 line 18-page 15 line 5; page 15 line 14-page 16 line 2; page 16 line 17-page 17 line 7.
- Providing reports: page 15 line 14-page 16 line 2 and page 16 line 17-page 17 line 7;

- Satisfying expansion request by creating an amount of necessary storage before changing the size that will be obtained by the management module in reports from the source virtual disk: page 12 lines 7-20; page 15 line 14-page 16 line 2; page 16 line 17-page 17 line 7; Fig. 5 ref. 530.
- Changing reported size: page 15 line 14-page 16 line 2.

D. Claim 16

Claim 16 is a claim for an apparatus for dynamically resizing, in this case shrinking, mirrored virtual disks in a RAID storage system, comprising:

> first means for providing an interface to a storage system;
>
> second means for providing communication with host devices; and
>
> means, coupled to the host side interface and the storage system interface, for
>
>> receiving a request to dynamically resize mirrored virtual disks in a RAID storage system,
>>
>> manipulating RAIDs in the RAID storage system assigned to the mirrored virtual disks, wherein the means for manipulating further comprises means for specifying a size of a virtual disk and mapping the size of the virtual disk directly to all components of a mirror set, detaching any RAIDs that extend beyond the specified size of the virtual disk and truncating RAIDs to free up any excess physical segments back into the RAID storage system, and
>>
>> before the step of manipulating RAIDs, resizing the mirrored virtual disks, and providing the resized mirrored virtual disks for operation.

Claim 16 is an independent apparatus claim, portions of which are stated in means plus function form, that relates to a particular embodiment of dynamic resizing.

Aspects of the above claim are supported at least by the following portions of the specification and drawings:

- first means for providing an interface to a storage system: Fig. 1 ref. 122 and page 13 lines 6-9 (storage area network); Fig. 1 ref. 124 and page 13 lines 6-9 (point-to-point connection); Fig. 1 ref. 140 and page 13 lines 10-14 (network connection), and Fig. 1 ref. 134 and page 13 lines 15-20 (virtual representation of disk); Fig. 2 ref. 200 and 240 and page 13 lines 10-14 (management module connected to storage).

- second means for providing communication with host devices: Fig. 1 ref. 120, page 13 lines 1-5 and page 13 lines 6-9 (network, Ethernet); Fig. 1 ref. 122 and page 13 lines 6-9 (storage area network); Fig. 1 ref. 124 and page 13 lines 6-9 (point-to-point connection); Fig. 1 ref. 140 and page 13 lines 10-14 (network connection); and Fig. 1 ref. 134 and (virtual representation); Fig. 2 ref. 210 and page 14 lines 1-7 (bus interface); Fig. 2 ref. 214 and page 14 lines 8-13 (Fibre Channel Arbitrated Loop); Fig. 2 ref. 216 and page 14 lines 1-7 (point-to-point connection), Fig. 2 ref. 222 and page 14 lines 8-13 (Fibre Channel), page 14 lines 1-7 (switched fabric).

- means for receiving a request to dynamically resize and manipulate RAIDs: Fig. 2 ref. 200 and page 14 line 18-page 15 line 5 (management module); Fig. 2 ref. 267 (processor); Fig. 2 ref. 292 (memory); Fig. 2 ref. 268 (program storage device); and Fig. 2 ref. 290 (computer program); also, any of first means or second means listed above.

- Process for manipulating RAIDs: Fig. 4 and 6 and associated text.

- Before the step of manipulating RAIDs, resizing the mirrored virtual disks, and providing the resized mirrored virtual disks for operation: Fig. 6, page 8 lines 3-9, page 17 lines 8-16.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

In the final Office Action, all claims were rejected based on 35 U.S.C. § 103(a) as obvious. Specifically, claims 1, 2, 9, and 19-24 were rejected as unpatentable over Lubbers et al. (U.S. Patent 6,880,052), hereinafter *Lubbers* in view of Bridge (U.S. Patent 6,530,035), herein after *Bridge*. Claims 3, 17, and 18 were rejected as unpatentable over *Lubbers* in view of *Bridge* and Cabrera et al. (U.S. Patent 6,629,202), hereinafter *Cabrera*. Claim 16 was rejected as being unpatentable over *Lubbers* in view of *Bridge* and DeKoning (U.S. Patent 6,275,898), hereinafter *DeKoning*. Claim 16 was also rejected based on § 112 as failing to comply with the written description requirement.

VII.   ARGUMENT

A. § 103(a): Obviousness over *Lubbers* in view of *Bridge* (Claims 1, 2, 9, & 19-24)

1. Claim 1

Claim 1 describes a method to dynamically expand mirrored virtual disks in a virtual disk storage system. The mirrored virtual disks include a source virtual disk and at least one destination virtual disk. The source virtual disk receives a request to dynamically expand the mirrored set. A critical aspect of the claim is that the sizes of the destination virtual disks be increased *before* a new size of the source virtual disk is reported. It is the order in which steps are performed that allow the expansion of a mirroring pair to be done dynamically, without having to quiesce the system, the mirroring status of the pair being able to continue during and beyond the expansion.

Lubbers is the primary reference relied upon by the Examiner in rejecting Claim 1. The Examiner states that *Lubbers* "does not *expressly* disclose that [increasing respective sizes of each of the at least one destination virtual disk is] implemented before reporting a new storage size of the source virtual disk, and reporting the new size of the source virtual disk." OA at 4 (emphasis added). In other words, *Lubbers* fails to teach this most critical aspect of ordering when doing expansion of a mirrored pair.

The word "expressly" in the sentence above suggests, however, that one might infer Applicant's ordering from what *Lubbers* teaches. That is definitely not the case. With respect to a logical unit (LUN) of storage in a mirroring relationship, *Lubbers* states that "once a LUN is increased in size, the increase can be propagated automatically to other members of a copy set." Col. 4 lines 51-54 (emphasis added). It is easy to imagine many ways that automatic propagation can be done, so we cannot conclude which approach *Lubbers* uses, or even whether the propagation is done correctly.

As an example of the danger in applying the *Lubbers* approach dynamically, without quiescing the system, suppose that the source and destination virtual disks in a

-14-

mirrored pair A and B are already filled with data, and that their sizes are to be doubled in capacity from 1 to 2 TB. Suppose A has already been increased in size, and this size increase is to be propagated to B. Now suppose more data is being sent to the expansion portion of A before B has physical storage to accommodate the increased capacity. How is that handled so that B continues to mirror A accurately during and beyond the size increase? *Lubbers* does not say specifically, but there are strong hints from the handling of similar tasks. For example, "destination disks or replicas are created 'on-the-fly' in a manner that enables destination disk creation, allocation, resizing, and reconfiguration *with little interruption of operational data transaction.*" Col. 4 lines 61-64. Also, "To establish a connection with the newly created destination virtual disk, *the source disk is briefly quiesced* in operation 613, during which time operational data traffic with hosts 102 may be cached." Col. 12 lines 42-45 (emphasis added). Applicant's dynamical expansion approach, which relies on appropriate ordering of operations, does not require quiescing the source disk to maintain synchronization (i.e., ongoing mirroring) of the information on the source and destination disks through and beyond the expansion.

Quiescing and caching may be inadequate to handle time critical operations. Indeed, the quiescing approach is exactly what Applicant's "specific logical ordering" (Specification page 7 lines 6-9) is designed to address. As stated in the specification, "The most common approach today is simply to break mirrors prior to resizing and then re-establish the mirrors afterwards (inducing long periods of re-copying to achieve a mirrored state, during which mirror backups don't exist and inherently can put the customers data at undue risk)." Specification page 6 lines 14-17.

The Examiner cites *Bridge* as adding the reporting and sequencing lacking in *Lubbers*:

> Bridge discloses ... expanding or shrinking logical volumes by adding or removing extents wherein when the logical volume is configured to a new size, the new size is reported in

> logical volume directory; thus allowing I/O operations are allowed [*sic.*] on the logical volume (col. 16, line 32-col. 17, line 4; col.20, lines 1-33) wherein the added or removed extents may be mirrored (col. 17, line 5-col. 18, line 58)."

OA at 4. Significantly, like *Lubbers*, *Bridge* does not teach the correct ordering of operations when expanding a mirrored set of virtual disks. The first section of text cited above (col. 16, line 32-col. 17, line 4) deals with expanding a *single* logical volume, not a mirroring pair or set. The second cited section (col.20, lines 1-33) relates to shrinking, so it is irrelevant to this claim. The third cited section (col. 17, line 5-col. 18, line 58) pertains to allocating a *new* mirrored extent set. In other words, none of the cited sections of *Bridge* deal with the missing requirement, explicit in Applicant's claim, of resizing the destination virtual disks before reporting a new size of the source virtual disk.

Because the Examine discusses this first section of text cited from *Bridge*, entitled "Expand Logical Volume," at some length (OA p. 22 line 13 – p. 23 line 9), we will examine what it teaches if used to expand a pair of mirroring virtual disks, A and B. Applying the technique to A, we would in steps 2 and 3 add needed data extents and pointers, then in step 4 update the logical volume size in the logical volume directory entry, hence effectively reporting the change. Now the technique would be applied to B. In other words, *Bridge*: add extents to A, report A's new size, add extents to B, and finally report B's size; Applicant: add extents to both A and B, then have A report the size. These two approaches are very different.

The reason given by the Examiner for combining the references is that

> At the time of the invention it would have been obvious to a person of ordinary skill in the art to modify the system of Lubbers which provides source and destination virtual disks in a copy set and resizes these virtual disks in an automatic fashion wherein any changes to a source virtual disk are propagated to the destination virtual disk and further explicitly expand the size or perform changes of size of the copy set of source and destination virtual disks of Lubbers and later reporting the size of the copy set or source and destination virtual disks in the same manner that Bridge first resizes a logical unit and later reports the changes to the logical unit

> by updating directory tables in order to allow I/O access to the virtual disks, since Bridge
> discloses this provides the advantage of dynamically accommodating to system requirement
> changes in a mirrored system configuration (col. 3, line 45-col. 4, line 59; fig. 9 and related
> text).

OA at 5. But since neither of the references deal with the reporting sequence for expansion of virtual disks that are already mirrored, combining the two references for whatever reason simply does not yield Applicant's invention. (Applicant is not attacking the references individually, as indicated by the OA on page 24. If neither reference teaches an element/limitation of Applicant's claim, then neither does their combination.)

Applicant's ordering of reporting relative to resizing of the destination virtual disk(s) allows the expansion to truly be done dynamically. Information on the destination disk(s) can remain synchronized with the source without quiescing any of the disks. As discussed in the example above, under the *Lubbers* teaching of increasing the size of the source disk and then propagating the size increase to the destination disks, it is easy to see how the destination virtual disk might no longer reflect the source unless all disks are quiesced during the process of increasing the size; based on comments in related contexts, quiescing is the approach likely followed by *Lubbers*. In reply to this argument, made in substance previously by Applicant, the Examiner stated in the OA:

> Regarding Applicant's remarks referring to the combination of Lubbers and Bridge not
> teaching resizing propagated synchronously to second virtual disk, it is noted that these
> features are not recited in the rejected claim(s). Although the claims are interpreted in light of
> the specification, limitations from the specification are not read into the claims.

OA at 21 (citation omitted). With this comment, Applicant respectfully disagrees. It is not necessary to include in a claim the expected advantages of the combination of elements and limitations presented there. Consider, for example, a claim for some mechanical device, such as a puncher for punching holes in paper. The claim does not need to discuss the advantages over prior art punchers, say with respect to slides for

-17-

positioning the paper correctly; the claim only needs to discuss the structure through a set of elements and limitations. The applicant is then free to explain in the specification, or in remarks, why the structure is advantageous over the prior art, which is relevant to obviousness. Indeed, MPEP § 2106 lists statements of intended use, "adapted to" and "adapted for" clauses as language that may raise a question as to the limiting effect of the language of the claim, and examiners typically ignore such language in assessing claim scope.

2. Claims 2 and 19

Claims 2 and 19 should be allowable as dependent on claim 1.

3. Claim 20

Claim 20 is an independent method claim that should be allowable for the same reasons as claim 1.

4. Claims 21 and 22

Claims 21 and 22 are dependent claims that depend on claim 20, and should be allowable for that reason.

5. Claim 23

Claim 23 was rejected by the Examiner for the same reasons as claim 1. The reasons given above for allowance of claim 1 also apply to claim 23, and are hereby incorporated by reference. In addition, the Examiner chose to ignore elements and limitations of the claim, as follows:

> a management module that includes a host side interface adapted to *[interpreted as intended use, See MPEP 2106 II-C]* communicating with host devices, through which the management module is adapted to *[interpreted as intended use, See MPEP 2106 II-C]* report the size of the mirrored virtual disks and to receive a request to expand the mirrored virtual disks, and a storage system interface *[interpreted as intended use, See MPEP 2106 II-C]* communicating with the virtual disks that is adapted to *[interpreted as intended use, See MPEP 2106 II-C]*

> requesting the source virtual disk to expand and *[interpreted as intended use, See MPEP 2106 II-C]* obtain reports of the size of the virtual disks from the source virtual disk; and logic adapted to *[interpreted as intended use, See MPEP 2106 II-C]* provide reports of the size of the source virtual disk to the management module through the storage system interface

OA at 7-8 (emphasis in original).

According to MPEP § 2111, "During patent examination the pending claims must be 'given their broadest reasonable interpretation consistent with the specification.'" As discussed above, all the elements and limitations ignored by the Examiner based on § 2106 are supported by the specification. Section 2106 itself states that

> The subject matter of a properly construed claim is defined by the terms that limit its scope. It is this subject matter that must be examined. As a general matter, the grammar and intended meaning of terms used in a claim will dictate whether the language limits the claim scope. Language that suggests or makes optional but does not require steps to be performed or does not limit a claim to a particular structure does not limit the scope of a claim or claim limitation. The following are examples of language that may raise a question as to the limiting effect of the language in a claim: ...

(Emphasis on "may" added.) In this case, the Examiner gave no rationale why the various elements and limitations were ignored. Obviously, Applicant intended these elements and limitations to have limiting effect; after all, they also appeared, in a different form, in other claims. According to § 2011, "The broadest reasonable interpretation of the claims must also be consistent with the interpretation that those skilled in the art would reach." A person skilled in the art, as relevant here, might be a computer engineer with an undergraduate degree with some knowledge of storage systems. *It defies imagination to believe that such a person would interpret claim 16 as truncated by the Examiner to have the same meaning as the claim as presented by Applicant.* So, in particular, Applicant contends that claim 23 should be interpreted as *clearly intended* by Applicant. In general, the spirit of §§ 2106 and 2011, taken together, demands that a rationale, grounded in English grammar and common sense , be provided for each instance of Examiner omission. Without a rationale from the Examiner, how can an applicant hope to respond to such deletions?

6. Claims 24 and 9

Claims 24 and 9 depend on claim 17 and should be allowable for the reasons stated above.

B. § 103(a): Obviousness over *Lubbers* in view of *Bridge & Cabrera* (Claims 3, 17, &18)

Cabrera is cited for teaching "specifying a size for the virtual disk and mapping the size of the virtual disk is performed by an operating system." *Cabrera* does not address the missing requirements, discussed above, of claim 1. Hence, claims 3, 17, and 18, which all depend on claim 1, should be allowable.

1. Dependent Claim 17

Claim 17 should be allowed for a reason independent of, and cumulative to, the reasons already stated for allowing claim 1. The method in claim 17 includes the step of "providing by the source virtual disk continuous availability for normal disk access operations between the step of receiving a request and the step of reporting the new storage size of the source virtual disk." The Examiner states that "Lubbers discloses the host can continuously write to source (col. 12, line 38-col. 13, line 15)." This is irrelevant to claim 17 for two reasons:

- Figure 6 of *Lubbers*, to which the cited language refers, pertains to creation of a new destination virtual disk (items 607, 609, 611, 613, 615, 617, 619), and not to resizing of a mirrored pair or set; and

- The process which Figure 6 describes includes quiescing the source virtual disk (item 617) and copying data in background (item 619)—this is anything but "continuous availability."

Cabrera et al. (U.S. Patent 6,629,202, hereinafter *Cabrera*) is referenced as disclosing "logical volumes and their plex are dynamically mapped and resized under the control of the operating system without system disruption." *Cabrera* does not disclose the

relative timing or resizing and reporting in a mirrored pair required by claim 1. Consequently, claims 17 should be allowed independently of claim 1.

2. Dependent Claim 18

Claim 18, which depends on claim 17, should be allowable for a reason independent of, and cumulative to, the reasons already stated for allowing claims 1 and 17. Claim 18 provides for continuous mirroring between the step of receiving a request and the step of reporting the new size of the source virtual disk. The portions of *Lubbers* and *Cabrera* cited by the Examiner as teaching continuous mirroring do not pertain to a resizing operation.

C. § 112: failure to comply with written description requirement (Claim 16)

Claim 16 is the only independent claim for dynamically shrinking, as opposed to expanding a mirroring set of virtual disks, that is being appealed. The OA states:

> Applicant's Specification does not appear to provide support for the newly amended limitations of ... "before the step of manipulating RAIDs, resizing ... providing ..." (claim 16). Applicant relies on paragraph 0049 to show support for this limitation; however, paragraph 0049 of Applicant's Specification recites "the process is reversed for dynamically shrinking mirrored virtual disks in a RAID storage system, with the exception that when downsizing you may need to shrink beyond the granularity that you expanded by ... to shrink ... you would need to first reduce the size of the VDisk and its mirrors, then remove the 50 MB raid in both the source and destination, finally truncate the two 100 MB raids into 75 MB raids", without providing any explanation of when the virtual disks are made available for operation or that the manipulating occurs after the virtual disks available for operation.

The paragraph (Specification page 16 lines 6-11) cited in the above passage states that the process of shrinking is reversed from that of expanding. This is reflected in the changed order steps from Fig. 5 (expanding) to Fig. 6 (shrinking), and text in the

specification describing the two figures. In particular, the step of specifying the virtual disk size occurs last for expanding, and first for shrinking.

D. § 103(a): obviousness under over *Lubbers* in view of *Bridge* & *DeKoning* (Claim 16)

Claim 16 (actually, with regard to obviousness, the OA deals with claim 4, but claim 16 was addressed on the same grounds) was rejected as obvious over the combination of *Lubbers*, *Bridge*, and *DeKoning*. *DeKoning* was cited in reference to aspects of claim 16 dealing with RAID storage systems, and is irrelevant to our argument. Otherwise, the grounds for rejection parallel those of claim 1, but in this case, as pertaining to shrinking mirrored virtual disks.

Specifically, *Lubbers* is cited as disclosing "resizing members of a copy set dynamically wherein any change made to one LUN member of a copy set is automatically propagated to the other members." OA at 12. *Lubbers* states, "when changes are made to a dependent attribute [such as size] of one member of a copy set, the change is made automatically to each other member of the copy set." The term "automatically" is ambiguous, and might simply mean under computer control. *Lubbers* does not specify the order for reporting a reduction in the size of the copy (mirrored) set. *Bridge* deals with shrinking a single logical volume, and so does not provide this critical missing piece. Applicant's disclosure is specific: "The process is reversed for dynamically shrinking mirrored virtual disks in a RAID storage system, with the exception that when downsizing, you may need to shrink beyond the granularity that you expanded by." Page 16 lines 6-11.

In summary, key to claim 16 is that these are mirrored virtual disks. When they are downsized, the downsizing must be reported in a specific order relative to the actual shrinking of the available storage for the virtual disks. *Bridge* describes a process for downsizing individual logical volumes. *Lubbers* is silent about the timing of changing the reported size of the virtual disks and shrinking individual volumes in a

mirrored set. The order is important. By immediately reporting the virtual disks in the mirrored pair as downsized, the storage system can continue without interruption in service. Whatever reductions in number or sizes of physical disks used to implement the virtualization scheme can be done afterwards. Since *Lubbers* is silent on how sequencing is handled with respect to downsizing mirrored virtual disks, and since *Lubbers* suggests that resizing must be handled not dynamically, but by quiescing disks, and since *Bridge* only deals with downsizing a single virtual disk without teaching ordering for downsizing mirrored pairs, the combination of *Lubbers* and *Bridge* does not yield Applicant's invention as defined by claim 16.

VIII. REQUEST FOR RELIEF

Based on the above rationale, the Applicant has appealed the Examiner's final rejection of the pending claims. The Applicant respectfully solicits the Board

Respectfully submitted,

Date: 11/24/2010

/James C. Evans/

James C. Evans

Registration No. 56,730

Beck & Tysver, P.L.L.C.

2900 Thomas Ave., #100

Minneapolis, MN 55416

Telephone: (612) 915-7006

Fax: (612) 915-9637

IX. CLAIMS APPENDIX

1. A program storage device readable by a computer embodying in a tangible medium one or more programs of instructions executable by the computer to perform a method for dynamically expanding mirrored virtual disks in a virtual disk storage system, the method comprising:

    receiving by a source virtual disk a request to dynamically expand the mirrored virtual disks, which include the source virtual disk and at least one destination virtual disk;

    associating additional storage with the mirrored virtual disks;

    increasing respective sizes of each of the at least one destination virtual disk before reporting a new storage size of the source virtual disk; and

    reporting the new size of the source virtual disk.

2. The program storage device of claim 1 wherein the step of associating additional storage further comprises:

    creating an amount of storage by providing RAIDs on each subsystem that is associated with each component of a mirror set;

    assigning the RAIDs to a specific virtual disk for a mirror device; and

    specifying a size for the virtual disk and mapping the size of the virtual disk directly to all components of the mirror set.

3. The program storage device of claim 2 wherein the specifying a size for the virtual disk and mapping the size of the virtual disk is performed by an operating system.

9. The apparatus of claim 23, wherein creating an amount of necessary storage includes providing RAIDs on each subsystem that is associated with each

-25-

component of a mirror set, attaching the RAIDs to a specific virtual disk for a mirror device and specifying a size for the virtual disk and mapping the size of the virtual disk directly to all components of the mirror set.

16. An apparatus for dynamically resizing mirrored virtual disks in a RAID storage system, comprising:

> first means for providing an interface to a storage system;
>
> second means for providing communication with host devices; and
>
> means, coupled to the host side interface and the storage system interface, for
>
>> receiving a request to dynamically resize mirrored virtual disks in a RAID storage system,
>>
>> manipulating RAIDs in the RAID storage system assigned to the mirrored virtual disks, wherein the means for manipulating further comprises means for specifying a size of a virtual disk and mapping the size of the virtual disk directly to all components of a mirror set, detaching any RAIDs that extend beyond the specified size of the virtual disk and truncating RAIDs to free up any excess physical segments back into the RAID storage system, and
>>
>> before the step of manipulating RAIDs, resizing the mirrored virtual disks, and providing the resized mirrored virtual disks for operation.

17. The program storage device of claim 1, the method further comprising:

> providing by the source virtual disk continuous availability for normal disk access operations between the step of receiving a request and the step of reporting the new size of the source virtual disk.

18. The program storage device of claim 17, the method further comprising:

  providing by the at least one destination virtual disk continuous mirroring of the source virtual disk between the step of receiving a request and the step of reporting the new size of the source virtual disk.

19. The program storage device of claim 1, wherein a first of the mirrored virtual disks has a different virtualization configuration from a second of the mirrored virtual disks.

20. A method, comprising:

  receiving a request to dynamically resize mirrored virtual disks, the mirrored virtual disks comprising a source virtual disk and a set of destination virtual disks that includes at least one destination virtual disk;

  associating additional storage with the mirrored virtual disks;

  increasing respective new storage sizes of each destination virtual disk before reporting a new storage size of the source virtual disk; and

  reporting the new storage size of the source virtual disk.

21. The method of claim 20, wherein the request is received by the source virtual disk.

22. The method of claim 21, wherein, in the step of receiving, the request is received electronically from a host and, in the step of reporting the new storage size of the source virtual disk, the new storage size of the source virtual disk is reported to the host.

23. An apparatus, comprising:

a set of mirrored virtual disks, including a source virtual disk and at least
one destination virtual disk, the at least one destination virtual disk
mirroring the source virtual disk, wherein the source and
destination virtual disks have the same size;

a management module that includes
a host side interface adapted to communicating with host devices,
through which the management module is adapted by logic
to report the size of the mirrored virtual disks and to receive a
request to expand the mirrored virtual disks, and
a storage system interface for communicating with the virtual disks
that is adapted to requesting the source virtual disk to
expand and to obtain reports of the size of the virtual disks
from the source virtual disk; and

logic adapted to
provide reports of the size of the source virtual disk to the
management module through the storage system interface,
satisfy an expansion request by creating an amount of necessary
storage before changing the size that will be obtained by the
management module in reports from the source virtual disk;
and
change the size that will be obtained by the management module
in reports from the source virtual disk.

24. The apparatus of claim 23, further comprising:
a host device adapted to send a request to the management module to
expand the mirrored virtual disks.

X. EVIDENCE APPENDIX

NONE

XI. RELATED PROCEEDINGS APPENDIX

NONE